Research on Intelligent Control Methodologies at the Instituto de Sistemas e Robótica of Instituto Superior Técnico*

Pedro Lima, João Sentieiro
Instituto de Sistemas e Robótica, Instituto Superior Técnico,
Av. Rovisco Pais, 1
1096 Lisboa Codex, PORTUGAL.
E-mail: pal@isr.ist.utl.pt

Abstract

The Instituto de Sistemas e Robótica (ISR) is a Portuguese research institution whose work focus on Systems Theory, Control, Robotics and Automation. In this paper we describe past and recent activity at the different Laboratories of the Instituto Superior Técnico pole of ISR, in the area of Intelligent Control. A particular approach to the design and evaluation of Intelligent Control Systems, proposed by ISR Intelligent Control Lab, is also summarized in the paper.

1 INTRODUCTION

The *Instituto de Sistemas e Robótica (ISR)* is a Portuguese research institution organized in three poles, located at Lisboa, Coimbra and Porto. The Lisboa pole is located within the campus of *Instituto Superior Técnico* (IST — Faculty of Engineering of the Technical University of Lisbon). Sixty seven researchers work at ISR/IST, including twenty seven who hold a Ph.D. degree.

In recent years, ISR/IST has been involved in a number of Portuguese and international R&D projects in the area of Robotics and Automation. The research interests of the Institute in this area are mainly focused on applications to land, underwater and aerospatial autonomous or semi-autonomous mobile platforms, robotic manipulators, and production planning and scheduling. Techniques used come from the fields of Automatic Control, Signal Processing, Artificial Intelligence, Computer Vision or Computer Science, to name a few.

The ISR/IST researchers currently working in Robotics and Automation have an historical background in Control, Signals and Systems theory. By the end of the 1980s there was a raising interest on the contributions from other fields, such as Artificial Intelligence and Computer Vision, to the Control and Operation of large-scale complex systems, ranging from chemical processes to robotic devices. The group participated actively in the European ESPRIT-II project AIMBURN ("Advanced Intelligent Multi-Sensor System for Control of Boilers and Furnaces") [1], whose objective was to improve the efficiency of a glass-melting furnace, based on optimal control techniques and the use of a vision system to extract features describing the flame and the non-melted materials inside the furnace [17]. The vision system provided additional information about the plant which was used by the control algorithms. A rule-based Real Time Expert Controller was also designed to coordinate the execution of the involved sub-systems and to handle exceptions, such as sensor/actuator faults or extreme situations incorrectly handled by the low-level controllers [8].

Intelligent Control is today a well-established field within the discipline of Control Systems. It emerged in the early 1970s as an attempt to characterize what distinguishes "machine intelligent" from "conventional" approaches, and to derive an analytic theory applicable, at least, to some classes of intelligent control problems. Since then, many different techniques have been used to provide alternative solutions for the control of dynamic systems, especially those whose mathematical modeling is difficult or even impossible. Fuzzy logic, artificial neural networks, genetic algorithms and expert systems are among those techniques. All those approaches share some biological background and aim at integrating conventional control theory with computer science and computer engineering in an efficient way, while keeping the traditional control engineering methodology of design[15, 16, 21].

The AIMBURN project was the starting point for the utilization of Intelligent Control techniques in an increasing number of projects at ISR. In 1995, two ISR researchers were co-organizers of the 1st *US/Portugal Workshop on Undersea Robotics and Intelligent Control* [14], held in Lisbon and joining about 50 American and European leading experts in the two areas. This paper is a survey of past and current research using Intelligent Control techniques, carried out at many of ISR/IST Laboratories.

The paper is organized as follows: in Section 2 we survey past and current contributions of researchers of the Institute, which was identified as being related to the area of Intelligent Control. Section 3 summarizes a methodology for the design and evaluation of intelligent controllers, proposed by ISR Intelligent Control Lab. Conclusions are reported in the final Section.

*Part of this research was supported by the Portuguese research program PRAXIS-XXI, under grant 3/3.1/TPR/23/94.

2 INTELLIGENT CONTROL AT ISR

Research which, according to the definition of the previous section, includes Intelligent Control techniques and/or aims at developing Intelligent Control methodologies, has been carried out in the last few years at the different ISR Laboratories. In this section we will summarize the most significant results accomplished so far. We have chosen to group the references by research areas, rather than by Laboratories where the work was actually developed.

Adaptive Control and Controller Design

Pioneer work on Adaptive Control focused mainly on the on-line modification of the control law, for closed loop systems whose controlled plant showed slowly time varying characteristics, due to changes on critical parameters and/or operating point (for smooth non-linearities). In recent years, many researchers have attempted to apply less "conventional" techniques, such as neural networks, fuzzy logic and genetic algorithms, to the adaptive control of strongly non-linear plants, or "difficult" linear systems, such as non-minimum phase systems, systems with delays, and systems difficult to model mathematically. Due to their time consuming nature, some of those techniques could only be applied so far to the off-line design of controllers.

At ISR, research on those topics started in the late 1980s. Initially, work was driven by the increasing interest on fuzzy logic controllers (FLCs). However, despite some interesting features, such as design simplicity and capability to incorporate human expertise, the lack of a design methodology driven by specifications for the controlled process, proved to be a major drawback of FLCs. Oliveira *et al*[12] approached this problem by introducing an adaptive fuzzy controller where a supervision level modified on-line an FLC operating in the control loop. The supervisor's algorithm consisted of changing the center of the membership functions describing the output linguistic terms of the controller, based on the difference between the observed and desired features of the closed loop response to step inputs. The features used were the overshoot and rise time. Supervisor action was triggered by the detection of a step in the loop set point. Some restrictions to the modification of the membership functions were implemented, to preserve the original linguistic meaning of the rules provided by a human expert. The same authors[11] presented later a slightly different approach, where the supervision was implemented by fuzzy meta-rules, and its application extended to the on-line adaptation of PID controllers. Oliveira *et al*[13] tackled the same problem using a different supervision technique — a genetic algorithm (GA). The parameters of a PID controller were adjusted in order to minimize the accumulated sum of the quadratic error between the loop set point and the plant output, for a given time horizon. The fitness function used by the GAs was the inverse of this performance index. The population of the GA was composed of strings concatenating a binary coded version of the three control parameters. Due to the non-predictable and time consuming nature of this technique, tests of tentative controllers on a real plant, before convergence to an acceptable solution, may be tedious and/or dangerous to the plant. Therefore, the methodology was suggested only for off-line controller design. Results comparable to those obtained by a Linear Quadratic controller designed for the same plant were obtained. Cardoso[3] proposed a GA-based approach to the supervision of FLCs, allowing the simultaneous determination of the set of fuzzy rules, the shape and location of membership functions and the universes of discourse of the linguistic variables. The performance index consisted of measures of the integral over time of the squared error between set point and plant output, or of the integral over time of the squared error multiplied by the time.

Vision-based Navigation and Guidance of Mobile Robots

One of the most powerful senses of human beings is the sight function. Thus, it is natural for mobile robots to use computer vision as one of its sensors, specially for navigation (self-location and/or obstacle avoidance) and guidance (path or track following) purposes.

A vision-based guidance controller for an Automated Guided Vehicle (AGV) was developed in the early 1990s by Cardoso *et al*[4]. The objective was to follow a track composed of a white strip on a dull black flat ground. The tracks considered had topologies frequently found in industrial plants, including straight lines, curves and crossing ways. The vision system acted as a smart sensor, capable of determining the error between the vehicle heading and the track tangent at each sampling instant, as well as predicting the track topology ahead of the vehicle. This information was used by a two-input two-output FLC to control the vehicle differential and common-mode speeds. Experimental tests made with this configuration showed the robustness of the guidance controller to noise in the track image, mainly due to dirty floors and shadows. The nature of the tracking problem also lead to modest computing times for image acquisition and processing. Finally, the FLC was relatively simple to design, and naturally prone to coping with different control states, such as trajectory junctions, derivations and crossings.

Most work on autonomous navigation of mobile robots using computer vision has been carried out in recent years at the Vision Laboratory of ISR. Santos-Victor *et al*[18, 19] described two *purposive* (in the sense of being specific to the problem) approaches to vision-based navigation. The first work uses the normal optical flow — computed over an image sequence acquired by a single camera — to detect obstacles lying on a flat ground. With this input information, the robot is capable of avoiding obstacles on the ground. Actually, holes in the ground can also be detected using this method, which dispenses any reconstruction of the environment surrounding the robot. At the initialization stage, the projective transformation between the image plane and the flat ground plane are estimated. Obstacles are detected by analysis of the inverse projection of the normal flow field onto the horizontal plane, noticing that (e.g., for translational motion) points on the ground plane should present the same flow vectors, while points lying above or below that plane will have respectively larger or smaller flow values. In the second work cited, a mobile robot equipped with a pair of cameras looking laterally and endowed with a controller based on the comparison between the apparent image velocity of the two cameras, is capable of navigating within narrow corridors with curves and lateral obstacles. One of the interesting features of this approach is the emulation of the behavior of freely flying honeybees. The navigation control loop is also studied analytically, based on

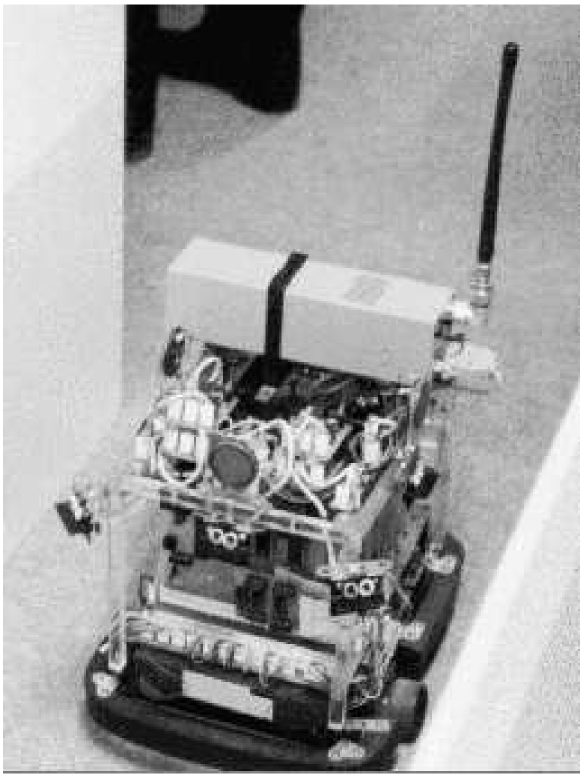

Figure 1: The MOONRAT vehicle.

control theory methods.

Small Autonomous Vehicles

Research on mobile robotics has always been driven by the ultimate goal of building intelligent autonomous machines. Recently, there is a growing scientific interest on the development of small, cheap and modular vehicles, endowed with distributed control systems and capable of integrating the information provided by several sensors, such as computer vision, sonar, encoders or gyros.

At ISR, work on small vehicles started in the 1980s, but it was usually directed towards specific applications, such as AGVs or "light followers". Later on, an important milestone was the MOONRAT vehicle (see Figure 1), a small ($37 \times 28 \times 34$ cm), light (4 Kg) semi-autonomous platform, equipped with rechargeable batteries, differential drive locomotion, eight infrared sensors for short distance and three sonar sensors for long distance obstacle detection, and underneath sensors to detect holes in the ground. A dedicated microprocessor handled all sensorial information, while a central microprocessor managed the flow of information between sensors and actuators (wheel DC motors), and implemented motor speed and vehicle guidance control algorithms. An user-friendly graphical interface was developed to help exchanging information between an external operator and the robot, which communicated through a radio link[23]. Recently, work on the development of small flexible AGVs has been mainly driven by a competition of mobile robots held annually at La Ferté Bernard, France — the *Festival International des Sciences et Technologies*. The participating robots must follow a 5 cm wide track painted on a flat floor, composed of straight lines and arcs (minimum radius = 1 m). The track and background colors alternate between black and white, in a chess-like pattern. Any track interruption must be detected, and the point where the track resumes after the interruption is signaled by a passive beacon located one meter above the ground. Billiard balls of three different colors (red, black and white) are scattered along the track and must be collected and discriminated. Only red balls should be transported until the arrival point. The first attempt to build such a robot was based on an originally teleoperated toy car. The vehicle had four driving/steering wheels, two driving DC motors and one steering servomotor, a closed loop motor speed controller (including an optical encoder), eight infrared (IR) emitter/receiver pairs to follow the track and one additional motor/IR system to search and move towards the passive beacon after a track interruption. A central microprocessor implemented the guidance control loop and coordinated the whole system[7]. Currently, an improved version of such a vehicle is being built (see Figure 2), with an original mechanical design which includes a step-motor/IR system to collect and discriminate the balls, computer vision to recover from track interruptions, and a distributed control system using a central 80486 motherboard and three PIC16C74 microprocessors.

Open Control Architecture for Robotic Manipulators

ISR has a PUMA 560 industrial manipulator which was used in the past with its original UNIMATE Mark-III controller, under VAL-II operating system. However, Mark-III controller architecture offers many obstacles to its use for "high level control", such as task planning and task coordinated execution. Despite its robustness and user-friendly features — such as programming in a high level language and the existence of a teach pendant — advanced control solutions, such as those based on vision and force sensors, are difficult to implement in this closed architecture. To overcome those problems, an open control architecture for the PUMA manipulator is currently under development at ISR. The methodology followed was the initial replacement of most control hardware by a board, manufactured by Trident Robotics, which allows direct access to the manipulator joint positions and torques by an external computer. Currently, a library of primitive tasks (e.g. move, plan trajectory, locate object) is being written to help future users of

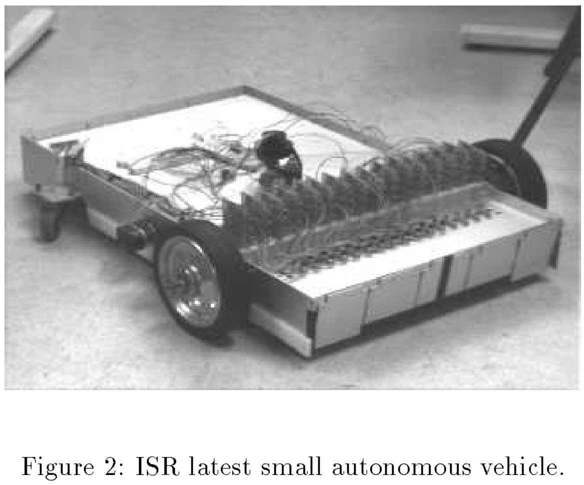

Figure 2: ISR latest small autonomous vehicle.

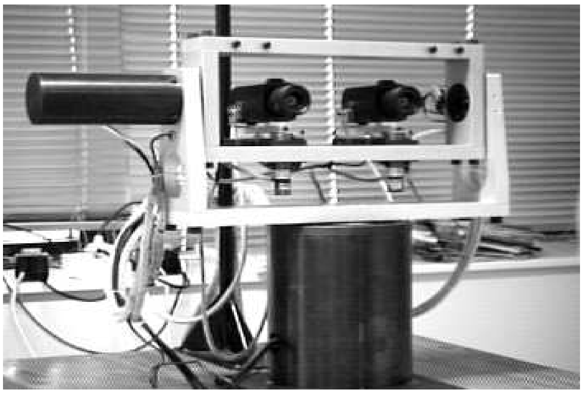

Figure 3: Medusa stereo head.

the testbed to develop their one task planners, force controllers or visual servoing algorithms, to name just a few. The control architecture is now open, in the sense that all involved algorithms are accessible to the programmer, including the parameter settings used. Joint-level trajectory planning and position control were already implemented on a Pentium PC (133 MHz) with a sampling frequency of 400 Hz and using different algorithms, with successful results[10]. Future work will include experiments on force control, visual servoing and coordinated execution of robotic tasks involving several pre-defined primitive tasks.

Cooperation Among Robotic Systems

The problem of coordinating multiple robot systems, with different types of motion constraints, either due to the mechanical design or the physical links between them, can be seen as a motion planning problem in an usually wide configuration space. Even though non-linear control techniques have been used to extend the analytic results obtained for systems with less degrees of freedom, this typically results into computationally heavy algorithms, and with little compliance to environment uncertainties. Sequeira *et al* [20] introduced a behavioral cooperation architecture for control of multiple robotic mechanisms. The architecture is composed of *behavior blocks, event sensor blocks* and a *task controller*. Each robotic device has an associated behavior block, which maps the configuration space of the device onto itself or one of its tangent spaces. Behavior functions inside each behavior block, are triggered by the occurrence of an event (or combination of events) detected by the event sensor block for the corresponding device. If an appropriate command for the robotic mechanism cannot be executed, another behavior function of the the same block is triggered. Events may result from environment sensing information or from state feedback from other robotic devices. For instance, if a manipulator mounted on a mobile platform cannot reach an object, it will create an event for the mobile robot device to move towards the object. The task controller is composed of a set of heuristic rules which define ranges of appropriate values for the timing of each task. The authors present results for a multiple robot system composed of a mobile platform, a 3 degree of freedom manipulator mounted on the platform, and a 3 degree of freedom gripper mounted at the tip of the manipulator.

Robotic Stereo Head

A robotic stereo head, named **Medusa** (see Figure 3), was designed and built at the ISR Vision Lab in the early 1990s. It has four mechanical degrees of freedom: common tilt, common pan and two independent vergences, and can be controlled by an external computer equipped with an image acquisition board. Some of the main ocular movements available in the human ocolumotor system, such as saccades, smooth pursuit and vergence, were implemented in an initial stage, at speeds and accelerations comparable to those of the human eye-system. In recent years, Medusa has been used to investigate the performance of several different visual servoing algorithms. Real-time vision-based control is frequently hard to accomplish, due to the long image acquisition and processing times involved. Speed increase is usually accomplished by a drastic reduction of the image resolution, thus reducing the accuracy of the control loop. Log-polar images provide an acceptable tradeoff for this problem. A log-polar mapping of a cartesian image reduces the overall amount of data, while keeping an acceptable resolution at the center of the image. Bernardino and Santos-Victor[2] described two different vergence control strategies, both based on the correlation of low resolution log-polar images. Besides being a especially suited technique for this control problem, log-polar mappings emulate the human retina, which contains a central area — the fovea — with a high density of photoreceptors, in contrast with its periphery. The authors investigated also the performance of different combinations of log-polar sensor layouts and control strategies, introducing a set of quality criteria for performance evaluation.

Production Planning and Scheduling

Production Planning and Scheduling (PPS) is a very important issue in the framework of Computer Integrated Manufacturing (CIM). As the management of production systems can be envisioned as a problem of controlling complex systems (non-linear and time-variant), the PPS field not only proposes rich research subjects, but also provides real-world testbeds for tools and approaches developed in different scientific areas.

Complexity in such systems also results: (i) from the high number of variables involved, and (ii) from the interaction among these variables. In such cases, the search for optimal solutions within analytical methodologies should be abandoned in favor of heuristic approaches. Artificial Intelligence — and particularly, expert systems, heuristic search, fuzzy logic, genetic algorithms, simulated annealing, among other research topics — have developed pragmatical approaches which are capable of handling the control of such systems.

The application of Intelligent Control techniques to Production Planning and Scheduling includes two main topics:

- *Fuzzy Logic in PPS* - Development and implementation of a methodology to control a hierarchical production system using Fuzzy Logic[5, 6];
- *Heuristic Search in PPS* - Application of heuristic search to production scheduling in a project job-shop problem[22].

Under the first topic, a hierarchical control structure based on three decision levels (higher, middle, lower), each responsible for a different production problem with a different time scale, is used. This methodology approaches the tasks associated with each level using a heuristic formulation and solves the short-range planning and scheduling problems with a nonstationary policy. The higher decision level determines safety stock levels used to compensate for future resource failures. At the middle level, loading rates are computed. This is accomplished through a fuzzy controller that tends to minimize the error between the cumulative production and the cumulative demand while keeping the work in progress below acceptable values. Finally, the lower level controls the flow of parts among resources, using a fuzzy decision method. This method has the ability to use several criteria to generate a decision. Simulation results reveal that this approach exhibits good performance, in terms of a high production percentage and a low WIP, even under resource failures and demand variations.

Regarding heuristic search in PPS, a project job-shop scheduling problem assumes that each job has a set of operations in a particular sequence, with the possibility of parallel execution of two or more operations. Also, associated with the job there is a starting date and a desired due date. The goal in a scheduling problem is to schedule the operations related to a set of jobs and determine which resources should be allocated to each operation, according to the technological constraints and due dates. The scheduling of a job-shop type production system is generally a NP-hard problem, so the computation time grows exponentially with the problem dimension. However, for most real-world problems, a pragmatic approach is usually followed: finding a reasonable solution in an acceptable amount of time. For this reason, the main goal of this research encompasses the efficient achievement of feasible solutions, satisfying technological restrictions, resource capacity constraints, and due dates regarding each job. The approach taken is based on a constraint management system allowing two types of solutions for project scheduling problems: quasi-optimal albeit time consuming, or reasonable ones in an acceptable computation time.

3 AN ANALYTIC APPROACH TO INTELLIGENT CONTROL

As mentioned before, there is an Intelligent Control Lab at ISR, whose activity started two years ago. The group running the Lab has an horizontal perspective, with the main purposes of applying Intelligent Control techniques to the control and operation of complex and/or large-scale systems, and of implementing control architectures encompassing low-level control, sub-systems coordination and task planning.

The theoretical analysis and synthesis of Intelligent Control systems is a key research topic for ISR Intelligent Control researchers. This section represents the point of view of the ISR Intelligent Control group on the future challenges faced by Intelligent Control as an autonomous discipline. A detailed description of these concepts can be found in Reference [9].

The major objective of an intelligent control loop is to accomplish a goal communicated by a command (a desired set point, in traditional control theory). However, intelligent machines often operate inside complex environments that disturb the expected results of their actions. These disturbances result from incomplete and/or imprecise environment modeling and unexpected events, such as hardware failures. Mathematical models of large complex systems are usually *incomplete* and/or *imprecise*. The reasons for this are either the difficulty to derive mathematical expressions fully describing the system behavior and/or the huge computational power needed to implement the model in a computer. One way to handle this problem is to model the *uncertainty* about the nominal system behavior instead. This is, roughly speaking, the approach of Stochastic Control and Robust Control to the modeling of control systems where decision making levels are not involved.

A performance index should then include some measure of the uncertainty on the system operation. However, the maximization of such an index would generate highly reliable controllers, but in most cases at the cost of heavy computational efforts. The proposed performance measure J is then a combination of *reliability* R and *cost* C,

$$J = 1 - R + C \ .$$

Reliability and cost must be defined in conjunction. Reliability is understood as the probability that a given algorithm will meet its specifications after a preset execution time. Cost — in terms of used computational resources of any kind — is measured for the worst sample of the information processed by the algorithm, i.e., for the sample that leads to the reliability closest to a pre-specified desired reliability. Notice that those definitions are general enough, so that they can be applied to different algorithms (e.g., control, image processing, computer vision, planning). Despite its generality, the definition forces the designer to pre-establish a performance measure for each of the algorithms composing the intelligent controller.

The performance measure can be used for off-line design of intelligent controllers, allowing a comparison among alternative designs. Nevertheless, it is especially suited for on-line improvement through feedback from the environment where the machine operates. In this case, the performance indexes of the different controller algorithms must be measurable on-line, and the controller must have a set of alternative algorithms for each primitive operation that must be implemented (e.g., a set of PID controllers with different gains, for a closed loop controller). A success/failure signal is generated after each execution of an algorithm, whether its specifications were met or not, respectively. This signal is used by a reinforcement learning algorithm to learn over time the best algorithm of the pre-defined set, given the current environment conditions. Lima and Saridis [9] proposed such an algorithm and proved its convergence with probability one for the best algorithm.

4 CONCLUSIONS

Intelligent Control has grown fast from an emerging to an established discipline within control theory. Its scope includes contributions not only from control theory, but also from other fields such as Computer Science and Computer Engineering. Such a diversity of contributions leads to a difficult definition of what is Intelligent Control. Still, a few common characteristics can be identified in the works of different authors: the emulation of biological systems, the use of non-conventional techniques, the attempt to formalize solutions which integrate conventional control with decision making, learning, computer engineering and computer science.

Work developed at the different Laboratories of the Instituto de Sistemas e Robótica of the Instituto

Superior Técnico of Lisboa clearly illustrates that diversity of contributions. This paper summarized research done at ISR/IST which shares some or most of the characteristics identifying modern Intelligent Control, from fuzzy adaptive controllers to a stereo vision head, passing by autonomous robot navigation and guidance, heuristic production planning and scheduling, open manipulator control architectures, small autonomous vehicles, robotic bees and cooperation between robotic systems.

Acknowledgments

The authors would like to thank their colleagues of ISR, Isabel Ribeiro, Carlos Pinto-Ferreira, José Santos-Victor, Luis Custódio, Paulo Oliveira and João Sequeira, not only for the tons of information they provided, but mainly for the work they carried out and/or supervised at the Institute, without which this paper would not be possible.

References

[1] AIMBURN. Final report of E.E.C. ESPRIT project 2192 — Advanced Intelligent Multi-Sensor System for Control of Boilers and Furnaces, 1992. M. Heitor, editor.

[2] A. Bernardino and J. Santos-Victor. Sensor geometry for dynamic vergence: Characterization and performance analysis. In *Proc. of Workshop on Performance Characteristics of Vision Algorithms — ECCV 96, Cambridge, U.K.*, 1996.

[3] F. Cardoso. Supervisão de controladores difusos: Aplicação de algoritmos genéticos e de procura no espaço de fase. Master's thesis, Instituto Superior Técnico, Av. Rovisco Pais, 1 — 1096 Lisboa Codex, PORTUGAL, 1996.

[4] F. Cardoso, F. Fontes, C. Pais, P. Oliveira, and M. I. Ribeiro. Fuzzy logic steering controller for a guided vehicle. In *Proceedings of MELECON 94*, 1994.

[5] L. Custódio, C. Bispo, and J. Sentieiro. A fuzzy decision system for production planning and scheduling. *IEEE Transactions on Robotics and Automation — special issue on Computer Integrated Manufacturing*, April 1994.

[6] L. Custódio and C. Pinto-Ferreira. Production planning and scheduling using a fuzzy decision system. In *Proceedings of Workshop on Fuzzy Scheduling Systems, 8th Austrian Artificial Intelligence Conference - Fuzzy Logic in Artificial Intelligence, FLAI 93*, 1993.

[7] D. Fernandes, L. Farrolas P. Brito, and P. Lima. Design and development of a flexible automated guided vehicle. In *Proc. of Teleman-HCM European Union Undergraduate Students Telerobotic Congress*, 1995.

[8] P. Lima, P. Oliveira, C. Ferreira, and J. Sentieiro. The architecture of a knowledge based controller for the operation of glass melting furnaces. In *Proceedings of the Workshop on Computer Software Structures Integrating AI/KBS Systems in Process Control*, may 1991.

[9] P. U. Lima and G. N. Saridis. *Design of Intelligent Control Systems Based on Hierarchical Stochastic Automata.* World Scientific Publishing, 1996.

[10] N. Moreira, P. Alvito, and P. Lima. First steps towards an open control architecture for a PUMA 560. In *Submitted to CONTROLO 96 Conference*, September 1996.

[11] P. Oliveira, P. Lima, and J. Sentieiro. Fuzzy supervision of direct controllers. In *Proceedings of the V IEEE International Symposium on Intelligent Control*, Sep 1990.

[12] P. Oliveira, P. Lima, J. Sentieiro, R. Bravo, R. Galan, and A. Jimenez. An architecture for the supervision of fuzzy controllers. In *Proceedings of the IROS'90*, July 1990.

[13] P. Oliveira, J. Sequeira, and J. Sentieiro. Selection of controller parameters using genetic algorithms. In *Proceedings of the European Robotics and Intelligent Systems Conference — EURISCON'91*, June 1991.

[14] A. Pascoal, P. Lima, F. L. Pereira, K. Valavanis, and G. Saridis, editors. *Proceedings of US/Portugal Joint Workshop on Undersea Robotics.* Instituto Superior Técnico and University of Southwestern Louisiana, 1995.

[15] K. Passino. Bridging the gap between conventional and intelligent control. *IEEE Control Systems Magazine*, 13(3):12–18, June 1993.

[16] K. Passino and U. Ozguner. Intelligent control: From theory to application. *IEEE Expert*, 11(2), April 1996. Special Track on Intelligent Control.

[17] J. Santos-Victor, J. P. Costeira, J. A. Tomé, and J. J. Sentieiro. A computer vision system for the characterization and classification of flames in glass furnaces. *IEEE Trans. on Industry Applications*, 29(3):470–479, May/June 1993.

[18] J. Santos-Victor and G. Sandini. Visual-based obstacle detection: a purposive approach using the normal flow. In U. Rembold et al, editor, *Intelligent Autonomous Systems.* IOS Press, 1995.

[19] J. Santos-Victor, G. Sandini, F. Curotto, and S. Garibaldi. Divergent stereo in autonomous navigation: From bees to robots. *International Journal of Computer Vision*, 14:159–177, 1995.

[20] J. Sequeira, M. I. Ribeiro, and J. Gonçalves. Cooperation among multiple robotic mechanisms. In *Proc. of SIRS96 Conference*, July 1996.

[21] Technical Committee on Intelligent Control. Report of task force on Intelligent Control, IEEE Control Systems Society. *IEEE Control Systems Magazine*, 14(3), June 1994. P. Antsaklis, editor.

[22] J. Viegas, L. Custódio, and C. Pinto-Ferreira. CROSS: Combined relation-operation based scheduling system. In *Proc. of the Fifth International Workshop on Project Management and Scheduling*, 1996.

[23] H. Zacarias, L. Moreira, R. Antunes, J. Sequeira, and M. Ribeiro. MOONRAT II — um robot móvel de pequenas dimensões. In *Proceedings of the CONTROLO 94 Conference*, June 1994. In Portuguese.